The SciDAC Visualization and Analytics Center for Enabling Technology

Program Review

Rockville, MD

4/28/2009

E. Wes Bethel, LBNL

Coordinating PI, co-PI

Overview/Objectives

- Introduce you to VACET mission.
- Convince you that we are meeting/exceeding expectations.
  - Define expectations, objectives
  - Present results that show excellent progress on all major program thrust areas.
- Solicit input for improvement.

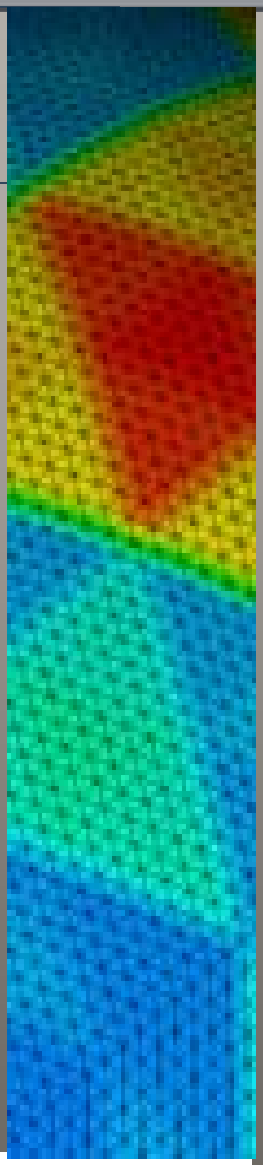

Talks Lineup

Overview, Executive Summary	Wes Bethel, LBNL
Science Impact	Sean Ahern, ORNL
Impact on Vis Community, VACET Research Portfolio	Ken Joy, UCD
Software Engineering	Hank Childs, LLNL
End-to-end science case study	Valerio Pascucci, Utah
Demo, future work, summary	Chris Johnson, Utah
Wrap-up	Wes Bethel, LBNL

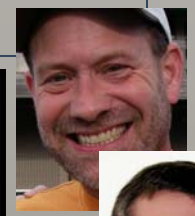
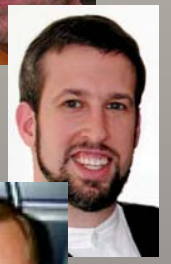
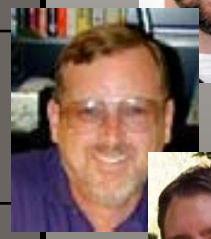
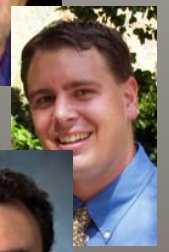

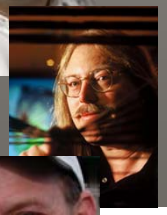


This Talk

- Part One: Problem Definition
- Part Two: Accomplishments Summary
- Part Three: Secret of our Success
- Part Four: Progress Against Proposal Milestones
- Part Five: The Road Ahead

Part One – Problem Definition

- SciDAC Program Goals
  - Enable science via use of HPC platforms.
- How?
  - Science applications
    - Produce new science (Climate, fusion, astro, etc.)
  - Centers
    - Provide production-quality, parallel capable s/w infrastructure for enabling science at the petascale, research component.
  - Institutes
    - Long-term research focus, outreach, training
  - Scientific Application Partnerships
    - “Code teams” of Science Apps, Centers, and others to transfer technology from Centers, Institutes into working SA code.

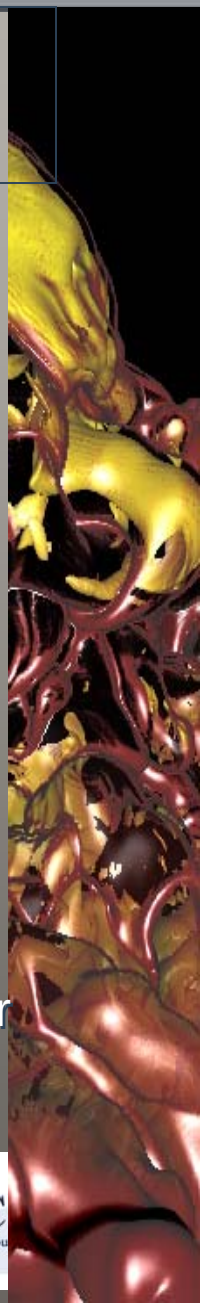

VACET Mission and Vision

- Mission: Leverage sci-vis and analytics software technology as an enabling technology for *enabling scientific insight*.
- Vision: adapt, extend, create, and deploy data understanding technologies for science stakeholders *to enable petascale science.*
- As a center, well positioned to respond to diverse needs/objectives through coordinated R&D, software engineering, outreach efforts.

How Enable Petascale Science?

- Fishing Analogy

Invent new fishing equipment	New algorithms, techniques, production-quality parallel capable s/w.
Teach others to fish	Partnerships with stakeholders, help transition communities to new technologies, outreach, training.
Catch big fish	Apply techniques to enable new science

Part Two - Accomplishments

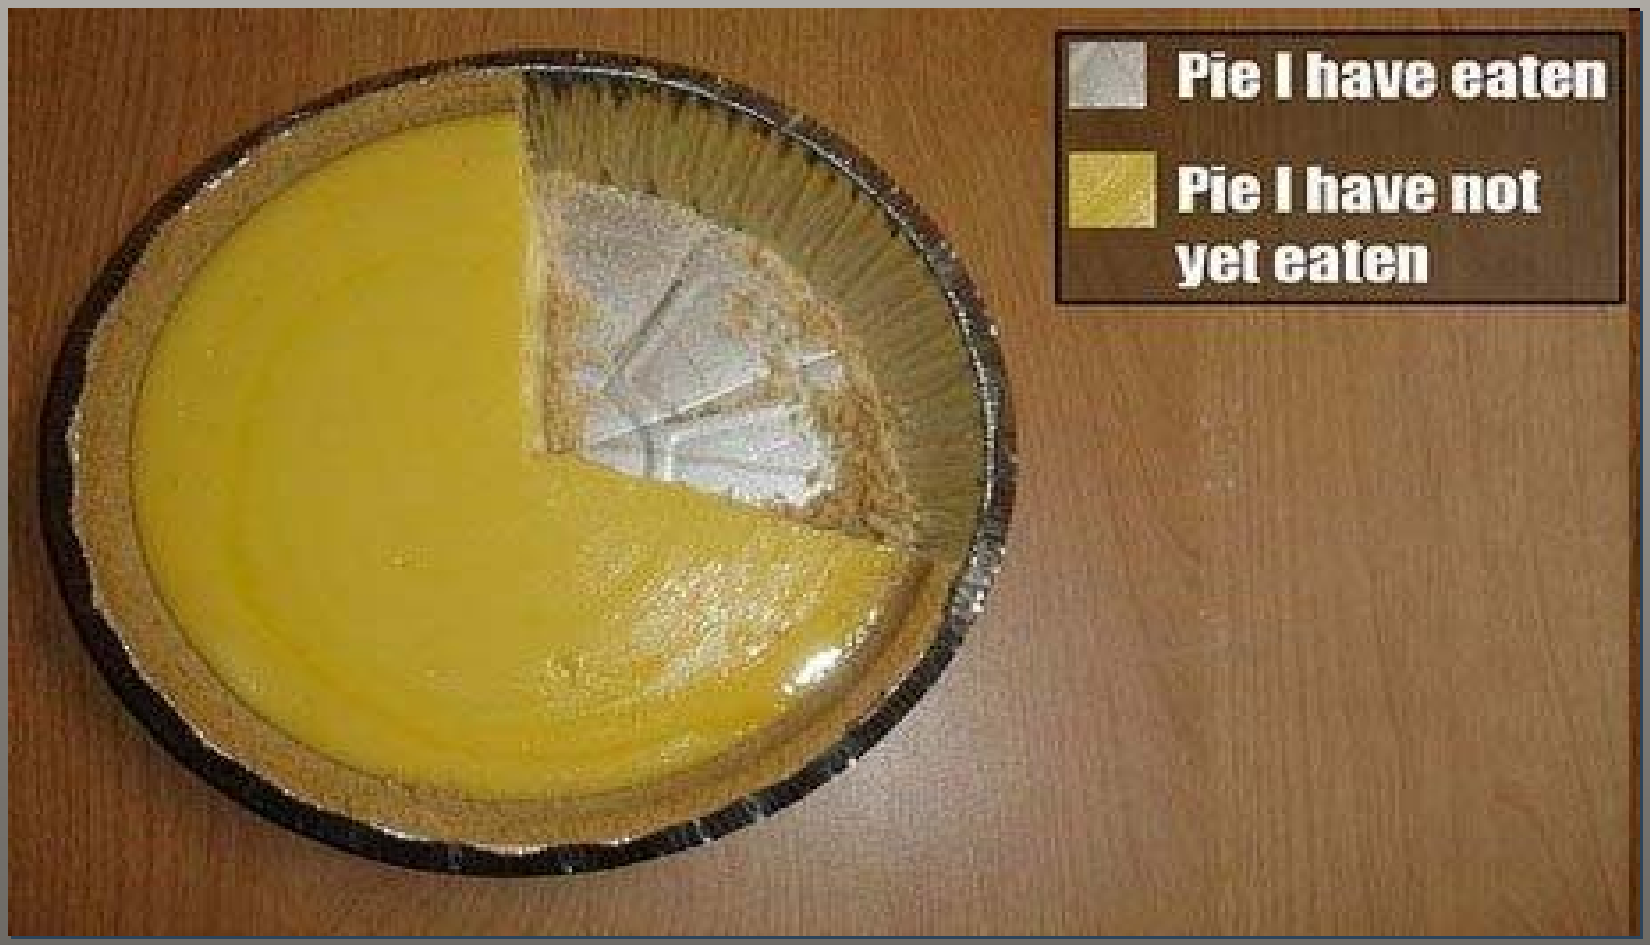

VACET Accomplishments Meet SciDAC Program Objectives

- Science Impact
- Production-Quality, Petascale capable software infrastructure
- Field-leading, award-winning research
- Progress towards petascale
- Effective use of ASCR Computing resources
- Service and Outreach

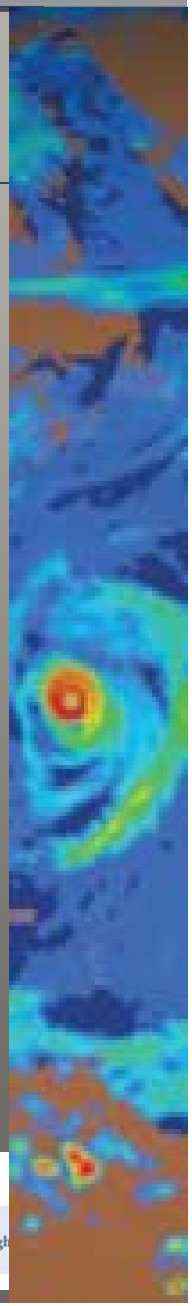

Science Impact (Sean's talk)

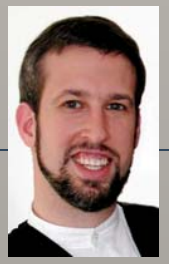

- New science insights.
  - New visualization R&D produces ability to quantify and display new phenomena never before seen (Valerio's talk: end-to-end science case study).
- Cost savings.
  - Communities and individuals adopt VACET technology allowing them to "buy rather than build."
- Improving efficiency of knowledge discovery.
  - Reduce time-to-solution, leverage parallel computing platforms (I/O, algorithms), delivering new capabilities in a familiar form.

Production Quality Software (Hank's talk)

- Objective: put production-quality, parallel capable visual data analysis technology in the hands of scientists.
- Results:
  - VACET technology installed an in use at all major DOE HPC facilities.
  - Truly "production quality."
  - Truly "parallel capable" on petascale platforms.
  - Effective business model for transitioning R&D into production software (CDAT, SLIVR, VisIt).

Field Leading Research (Ken's talk)

- Prolific, high quality publications:
  - ~70 journal articles, ~29 conference proceedings, ~6 invited articles, ~18 book chapters and edited books, ~10 posters.
    - All major visualization/graphics venues: IEEE Visualization, Siggraph, TVCG, Eurovis.
  - IEEE Visualization "Best Paper" or "Best Application Paper" 2006, 2007, 2008.
- VACET research portfolio driven by stakeholder needs. Techniques for:
  - Large data vis; scalable vis; flow vis; feature detection, analysis, and tracking; multivariate and temporal vis; statistical and uncertainty vis.

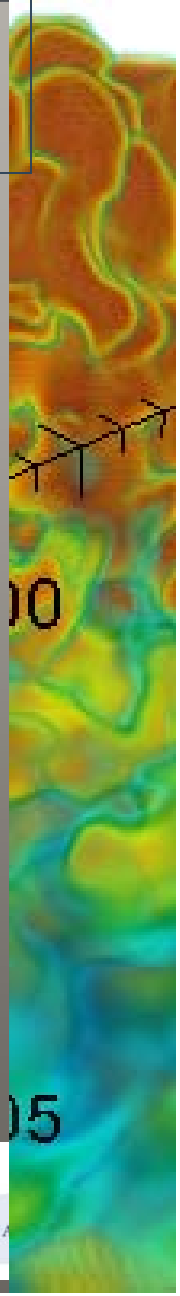

Progress Towards Petascale

- Help DOE show effective use of petascale platforms:
  - VisIt is the only visualization code that is part of the Joule metric. (Hank's talk)
  - Recent results: 1T zones on 40K jaguar cores!
- Effective tools and techniques for visual data analysis of data being produced by petascale applications on petascale platforms.
- Close coordination with staff at HPC centers to address end-to-end scientific stakeholder challenges (e.g., I/O).

Effective use of ASCR Computing Resources

- Deployment: VACET software installed and in use today at all major DOE HPC facilities.
- Synergy:
  - VACET team includes visualization and analytics staff from two major DOE HPC facilities.
  - Close working relationship between VACET and other staff (e.g., consulting, systems) at DOE's major HPC facilities.
- Resources
  - VACET has ERCAP allocation at NERSC. We use NERSC heavily for major visual data analysis projects (need more disk space and interactive analysis capabilities!)
  - VACET has Director's allocation at ORNL. Used for performance testing, scalability studies (e.g., Joule work).

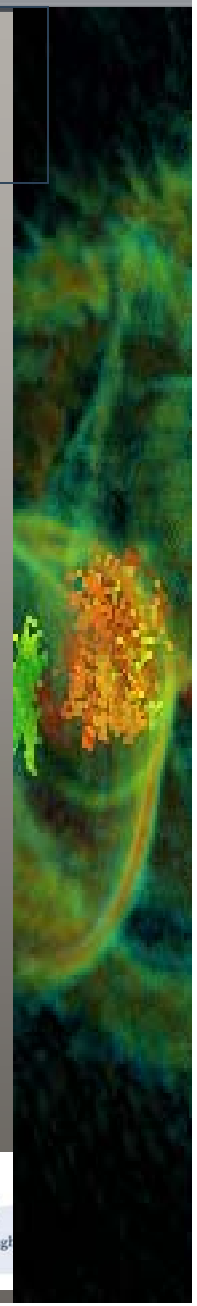

Service and Outreach

- Outreach to science stakeholders.
  - Tutorials: day-long at SciDAC program meeting, day-long onsite (e.g., PPPL, Fermi)
  - Workshops: Contributions to various IUSV workshops (e.g., SC 2007, 2008).
  - Speaking engagements: ~67 invited presentations.
- To the Visualization Community
  - Conference chair (9), program committee (60), technical reviewer (15*), advisory boards and national commitees (14).
- DOE-wide service
  - Numerous workshops (17), technical reviewer (e.g., SBIR program)

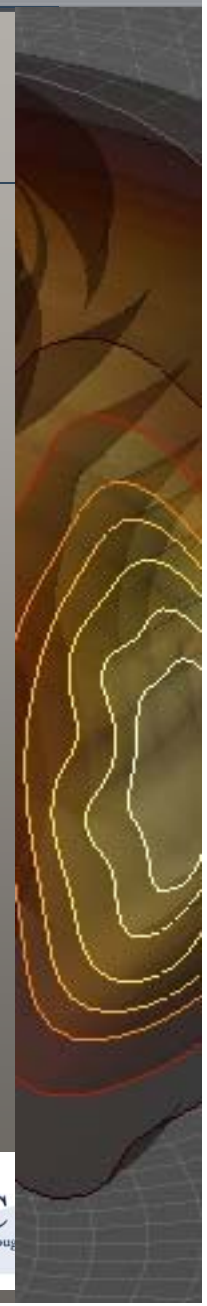

VisIt Tutorial at PPPL, September 2008

Congressman Jim Matheson (Utah), 4/2008

Awards – SciDAC 2008

- SciDAC 2008 "Viz Night"
  - Three "People's Choice" Awards
  - One Honorable Mention
  - Stakeholder(s) win People's Choice Awards using VACET s/w:
    - Tech-X: Accelerator
    - LLNL: NIF

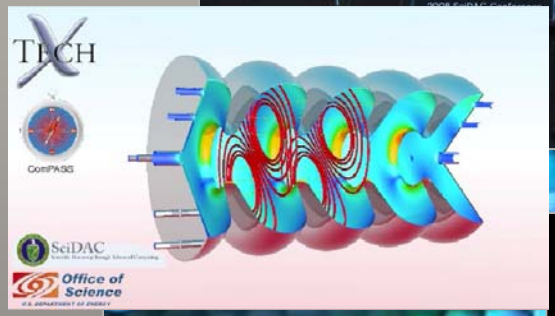

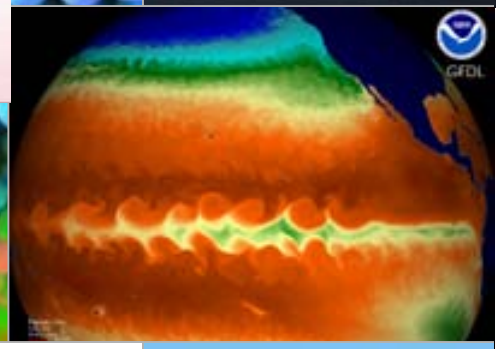

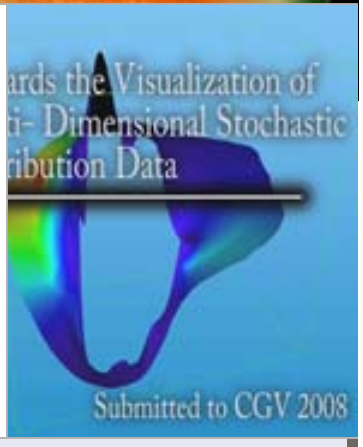

U.S. DEPARTMENT OF ENERGY | Office of Science

Collaborations/Stakeholders

- Who have we worked with?
  - Science stakeholders (more to come)
  - Technology providers (more to come)
  - Resource providers (more to come)
- Challenges
  - One-on-one projects not scalable to larger customer base
  - Generic software might not be responsive to stakeholder needs
- Response: focus on communities as middle ground.
  - Helping more than one person at a time
  - Well-focused needs drive R&D, which if successful, will meet needs of stakeholders.

Science Stakeholders/Collaborations

- *Accelerator* – reduce duty cycle in petascale data analysis.
- *Astrophysics* – community-wide infrastructure.
- *Climate* – new visualization capabilities delivered in production-quality form.
- *Combustion* – new science.
- *Fusion* – community-wide infrastructure, new high-performance capabilities.
- *Mathematics* – community-wide infrastructure.
- *Turbulence* – new visualization and high-performance capabilities.

Technology/Outreach Collaborations

- Mathematics – leveraging stakeholders, community-wide visualization infrastructure.
- Data Management – leverage index/query technology, tech transfer.
- Outreach Center – software engineering infrastructure, provide content for OC outreach activities.
- ITAPS – code interoperability.
- IUSV – joint work on projects, outreach.

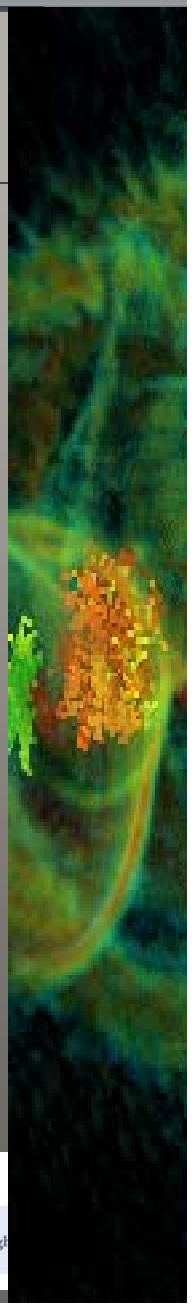

Other Collaborations

- DOE HPC Centers: NERSC/LBNL, NCCS/ORNL, ALCF/ANL:
  - Use of resources for R&D, application to stakeholder problems.
- NNSA – shared development, infrastructure maintenance for VisIt.
  - VACET: AMR visualization, query-driven visualization, open software engineering infrastructure, large data.

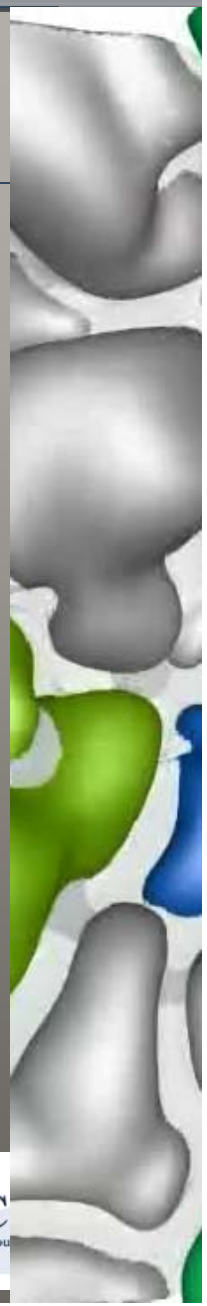

Accomplishments Summary

- Successfully brought multiple “products to market”
  - Science applications adopt VACET technology as community-wide visual data analysis s/w infrastructure.
  - Stakeholders are voting with their feet.
- Award-winning research
  - Dozens of peer-reviewed field-leading journal articles
  - Numerous Best-Paper awards
- High (and positive) visibility within the SciDAC and visualization communities.
- Realizing vision of a successful SciDAC Center

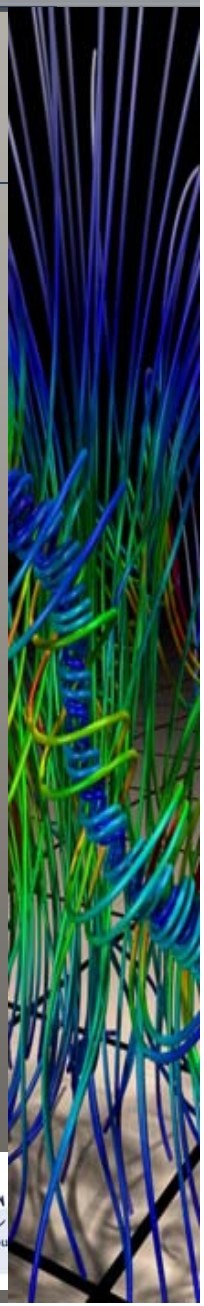

Part Three – The Secret of Our Success

- Team organization
- Customer focus
- Strategic partnerships
- VACET operations

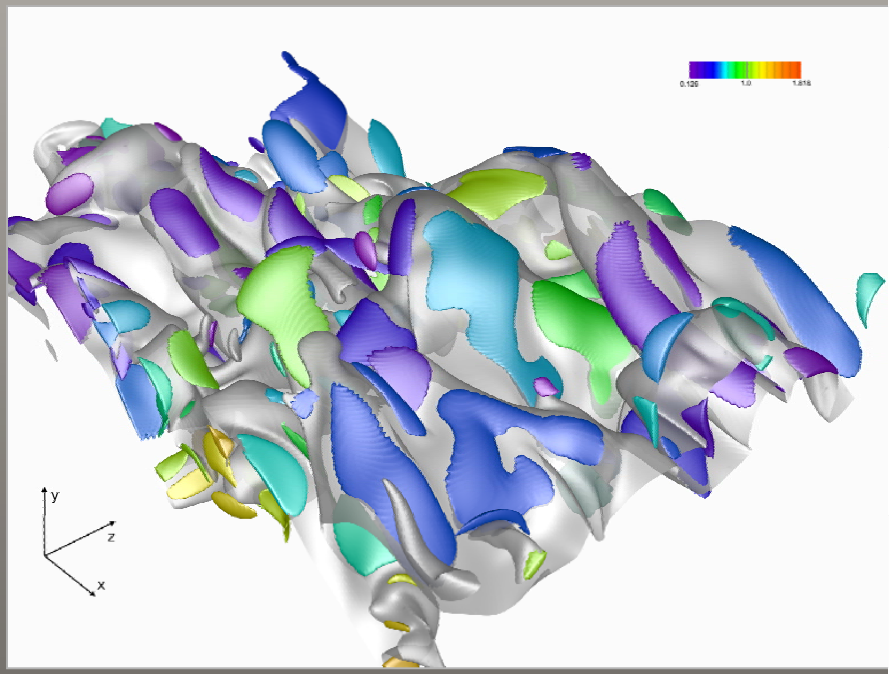

Team Organization

- Teams: stakeholder projects, R&D projects, software engineering projects.
- Executive committee: cross-institutional, cross-team coordination.

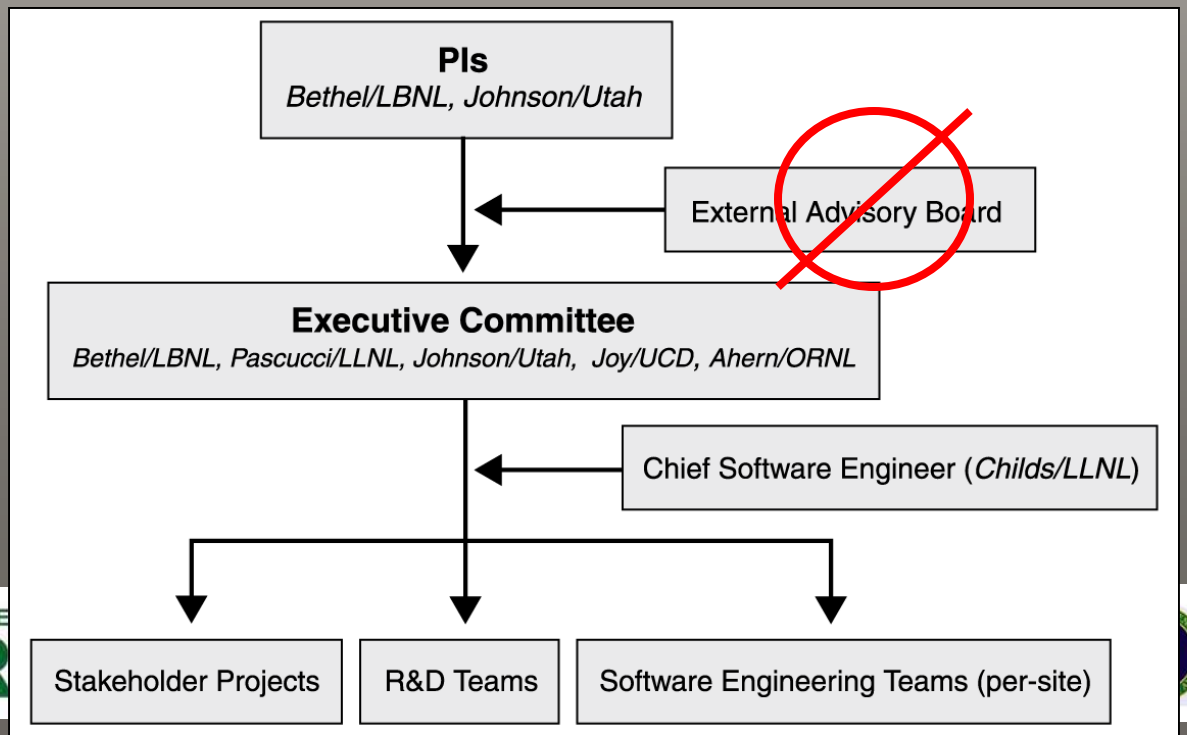

Customer Focus

- All R&D driven by stakeholder needs.
  - Results are virtually guaranteed to have positive scientific impact.
- Requires having customers ☺
- "Customer project manager" position in VACET plays a key role in helping to foster these relationships and regular interactions.
  - This position analogous to "product manager" in industry.

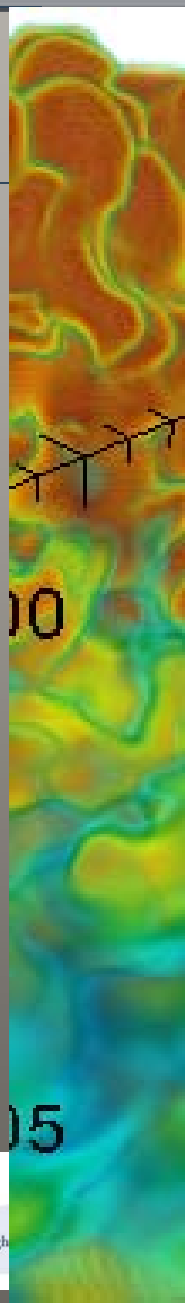

Strategic Partnerships

- Leveraging is key
  - Leverage stakeholders (APDEC story)
    - *Strong partnership that integrates Applied Math and CS with domain science.*
  - Leverage technology (FastBit and the SDM Center)
  - Leverage objectives: Outreach center mission
  - Leverage technology: resources
  - Needs common across multiple stakeholder groups allows for applying a technology to a larger group of stakeholders (accelerator, fusion, life sciences example).

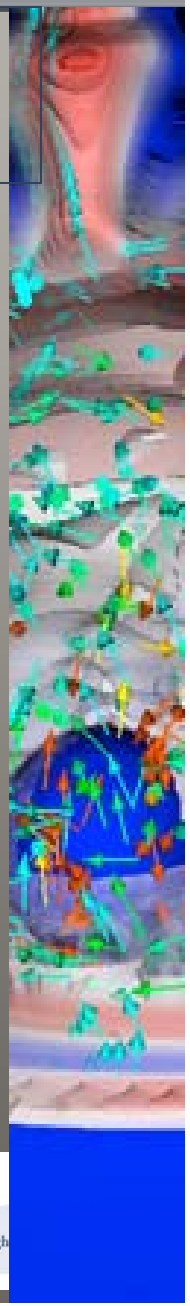

VACET Operations

- Regular EC communication: bi-weekly calls to coordinate activities, fine-tune objectives and direction.
- Twice-yearly All Hands Meetings
  - Present results, discuss priorities, look for opportunities to expand intra-team efforts.
- Website: www.vacet.org
- Twice-yearly progress reports to DOE.

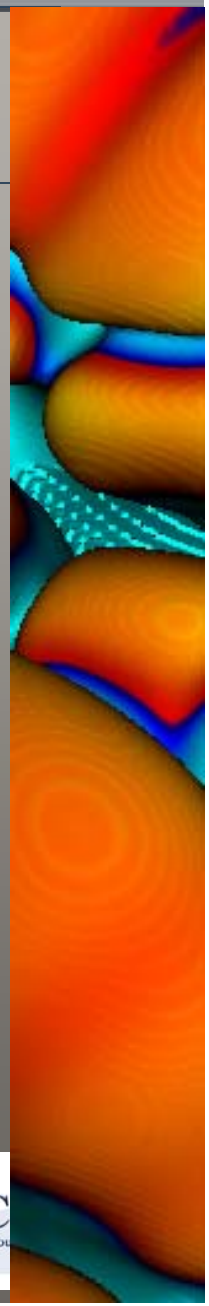

Project Management

- Twice-yearly internal review of all VACET activities.
  - Progress towards objectives/milestones
  - Revised milestones, cost estimate
  - EC reviews relative priority.
    - Some projects discontinued due to lack of funding, etc.
- This activity performed in months before the we compose the twice-yearly progress report for DOE.

Operations Case Study – APDEC

- Stakeholder project manager: G. Weber, LBNL
- Initial negotiations, list of high-level needs, priorities, programmatic targets:
  - Bethel, Colella
- Product development team:
  - Software engineering: Childs, Whitlock, ...
  - R&D team: AMR flow visualization: Garth, ...

Operations Case Study – LWFA/Accelerator

- Stakeholder Project manager: Bethel
- Initial negotiations:
  - Bethel, Geddes
- Product development team:
  - Visual data exploration: Prabhat, Rubel, Ahern, Childs, Meredith
  - Analysis: Ushizima, Rubel
  - Software engineering: Prabhat, Rubel, Childs, Meredith

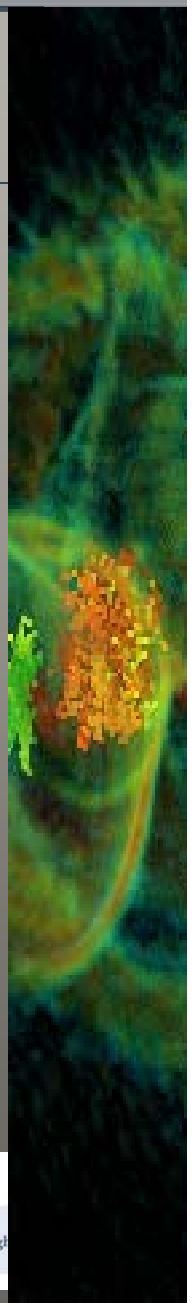

Who is doing What in the Project?

- Stakeholder project teams:
  - Primary stakeholder project manager
  - They rely on R&D teams, SWE teams, etc. to complete their mission
- R&D teams:
  - Primary team leader
  - Team includes others within VACET as needed
- Software engineering:
  - Primary software architect: Childs
  - Focused software engineering teams, project-oriented
- Management
  - Reports, interfacing with DOE, building/servicing external relationships, etc.

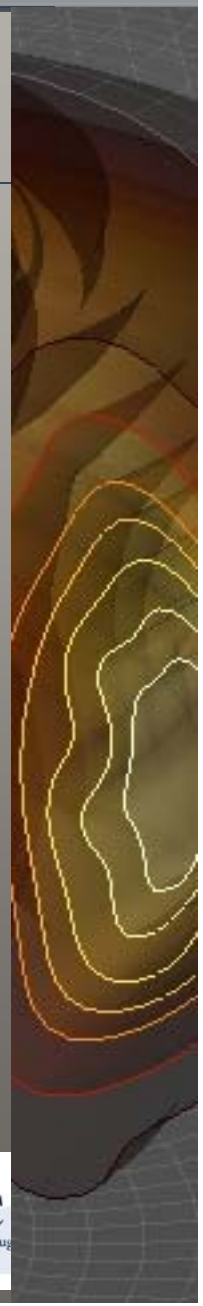

Part Four – Progress Towards Proposal Milestones

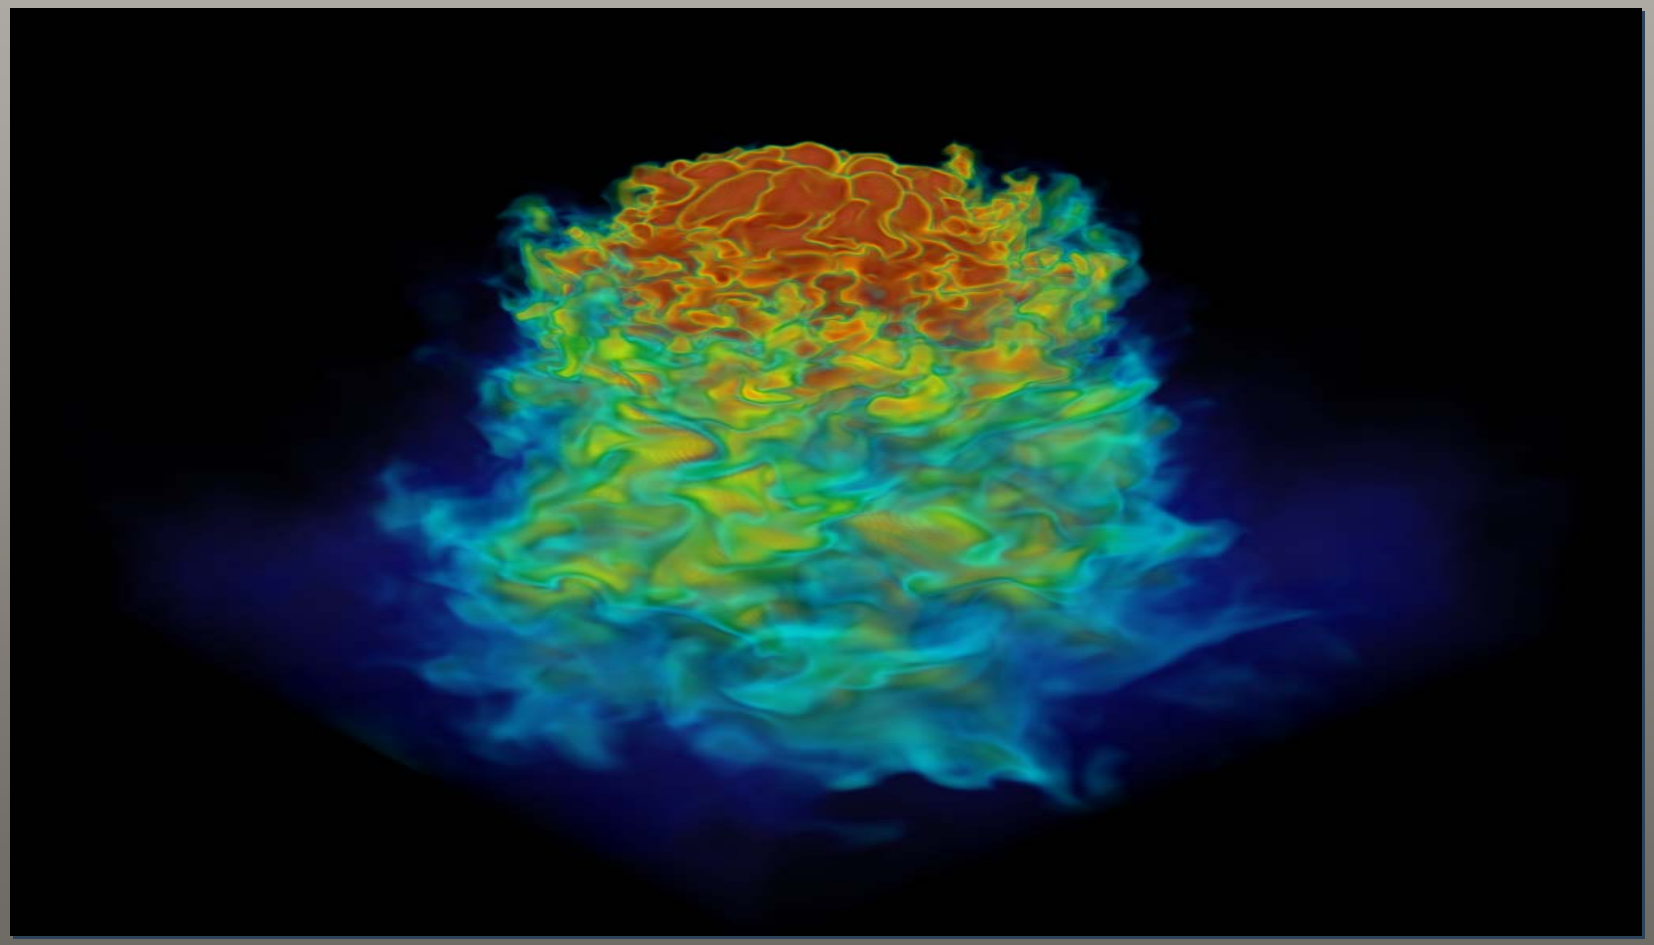

Progress Against Original Milestones

- High level view:
  - Achieving substantive impact on accelerator, astrophysics, combustion, fusion, mathematics, turbulence.
  - Performing excellent research in multiple directions aimed at petascale knowledge discovery: analysis, flow visualization, visual data exploration.
  - Effectively operating as a SciDAC center, including a high degree of leverage with related programs/strengths.

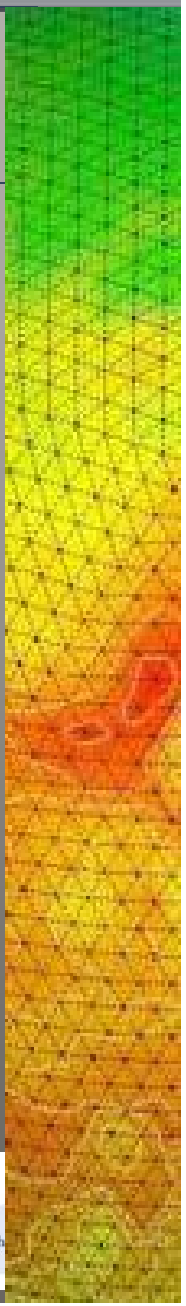

Scope Change Due to Budget Cut – High Level View

- Items specifically eliminated or not started:
  - Specific outreach plans called out in proposal.
    - VACET activities still include a lot of outreach.
  - Advisory board – travel expenses.
    - Our stakeholders are our advisory board – they give us direct feedback by telling us what is important to them.
  - Experimental science
    - Focusing primarily on computational science projects.
    - Flat budget requires releasing staff working on one such project.
  - Our Analytics effort is much more limited than planned – relatively few staff with analysis expertise, relatively few analysis projects.

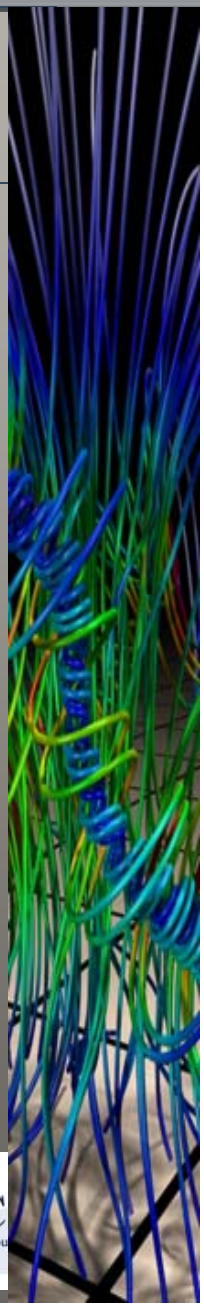

Scope Change Due to Budget Cut, ctd.

- Advanced analytics: a number of hoped-for projects likely will not be started.
  - Astro: feature mining: spectra to simulation model searches
  - Combustion: chemical pathway analysis/analytics; specialized “simulation dump” routines for streaming/multires; topological and semantic query-based feature definition and tracking; feature mining and visual correlations.
  - Climate: carbon cycle analysis, coupled model analytics
  - Fusion: topological and flow analysis
  - Turbulence: specialized “simulation dump” routines for streaming/multires; multires visualization/analysis; topological and statistical analysis; sim-sim and experiment-sim comparative analytics

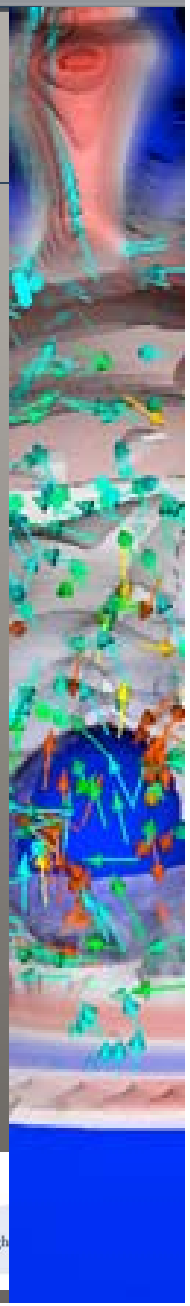

Scope Change Due to Budget Cut, ctd.

- Stakeholders we've not worked with due to insufficient funding:
  - Lattice QCD
  - Portions of Accelerator (EM, specifically)
  - Climate: Improving Global Climate Models
  - Fusion: Fusion Simulation Prototype projects (2006)
  - Groundwater/EM SciDAC projects.
  - Life Sciences SciDAC projects.
  - Materials/Chemistry SciDAC projects.

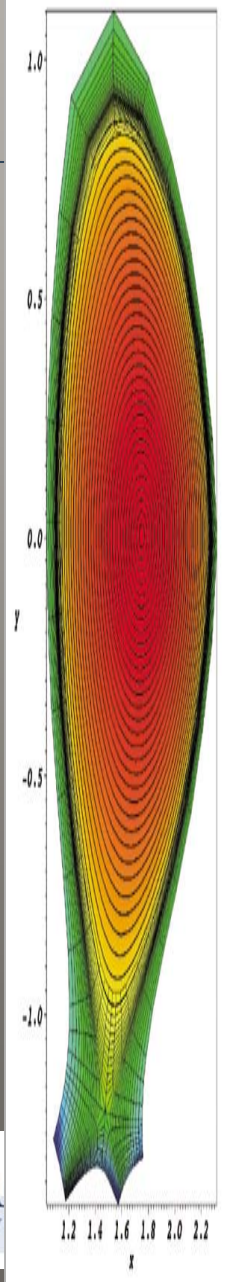

Relationship Between Publications and Work

- Publications taxonomy
  - "Applications" papers: stakeholder venues, visualization/analytics/computing venues.
  - Research papers:
    - New techniques that target short-term and long-term stakeholder capability requests.
    - New techniques targeting longer-term issues that need to be solved for effective petascale-class data visualization.
  - Invited articles:
    - What are you guys doing?
    - Please give us an expert opinion on <some topic>.
  - Posters
    - "Advertisements" for team's efforts.

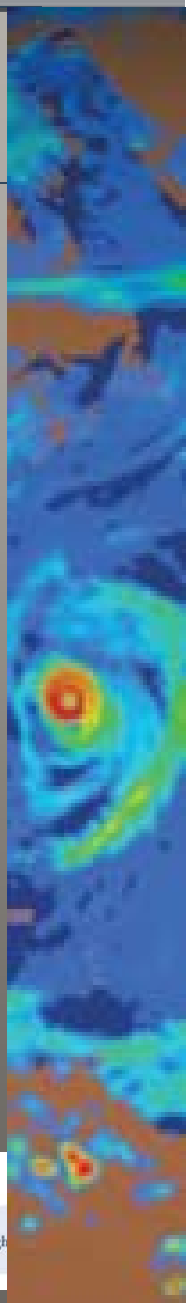

Part Five – The Road Ahead

Future Work – High Level View

- Stakeholder Projects
  - Many stakeholder projects underway, most “in the middle” of the “to-do” list.
  - These lists generally comprise a set of capabilities needed to realize vision of “enable petascale science” and have been sized to fit our available budget and personnel.

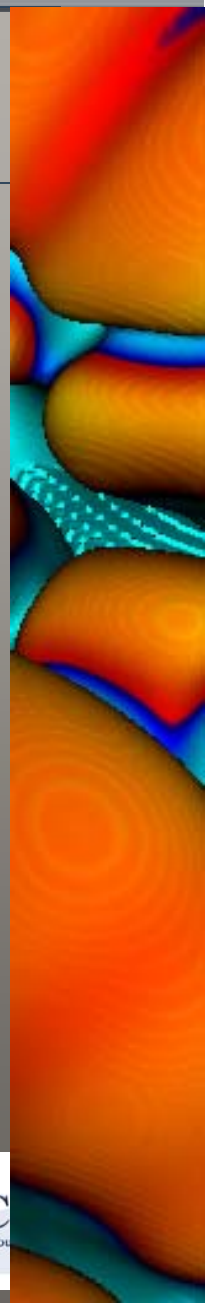

Future Work – In the Weeds: Accelerator

- Extend FastBit accelerations/QDV for use on field-based data.
- Implement “named selections” in VisIt
- Temporal parallel coordinates
- Larger datasets coming – will our s/w infrastructure work with them?
- Compute and display statistical information of fields associated with “interesting particles”
- Comparative vis/analysis of raw and summarized fields/particles
- Stable data model, parallel I/O
- Domain-specific vis: frequency maps
- Better temporal visualization/analysis

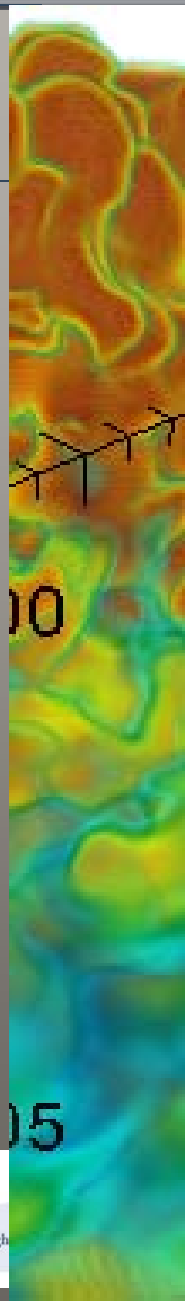

Future Work – In the Weeds

- Similar stories for:
  - Astrophysics
  - Climate
  - Combustion
  - Fusion

VACET Summary

- Producing positive scientific impact across many disciplines.
- Strong scientific community support.
- Award-winning research.
- Wildly successful, exemplary performance as a CET.

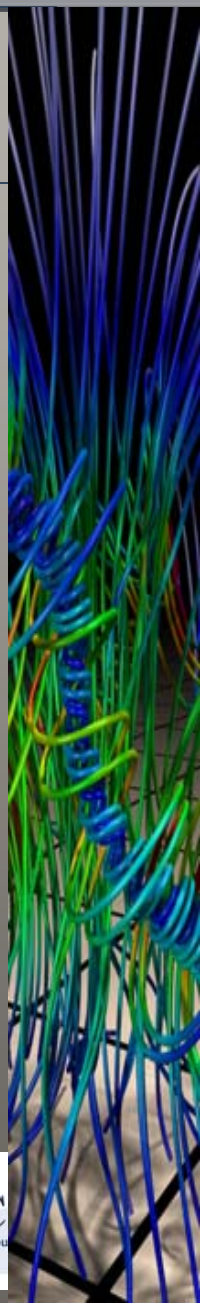

Parting Comments

- VACET has delivered on its promise.
  - Measureable science impact.
  - Progress towards petascale.
  - Field-leading research.
  - Effective use of ASCR computing resources.
- We are a highly productive, well integrated team that is making excellent progress.

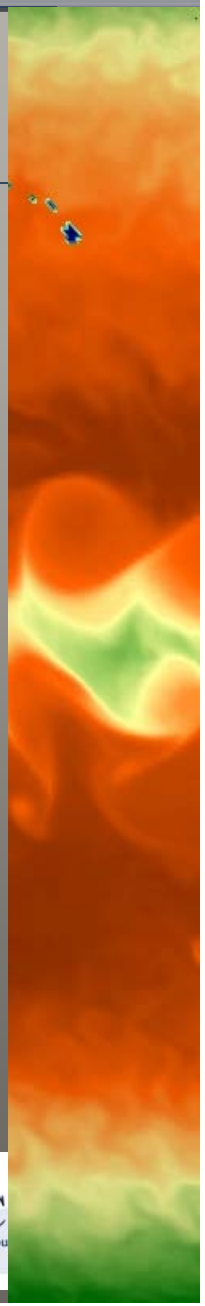

The End

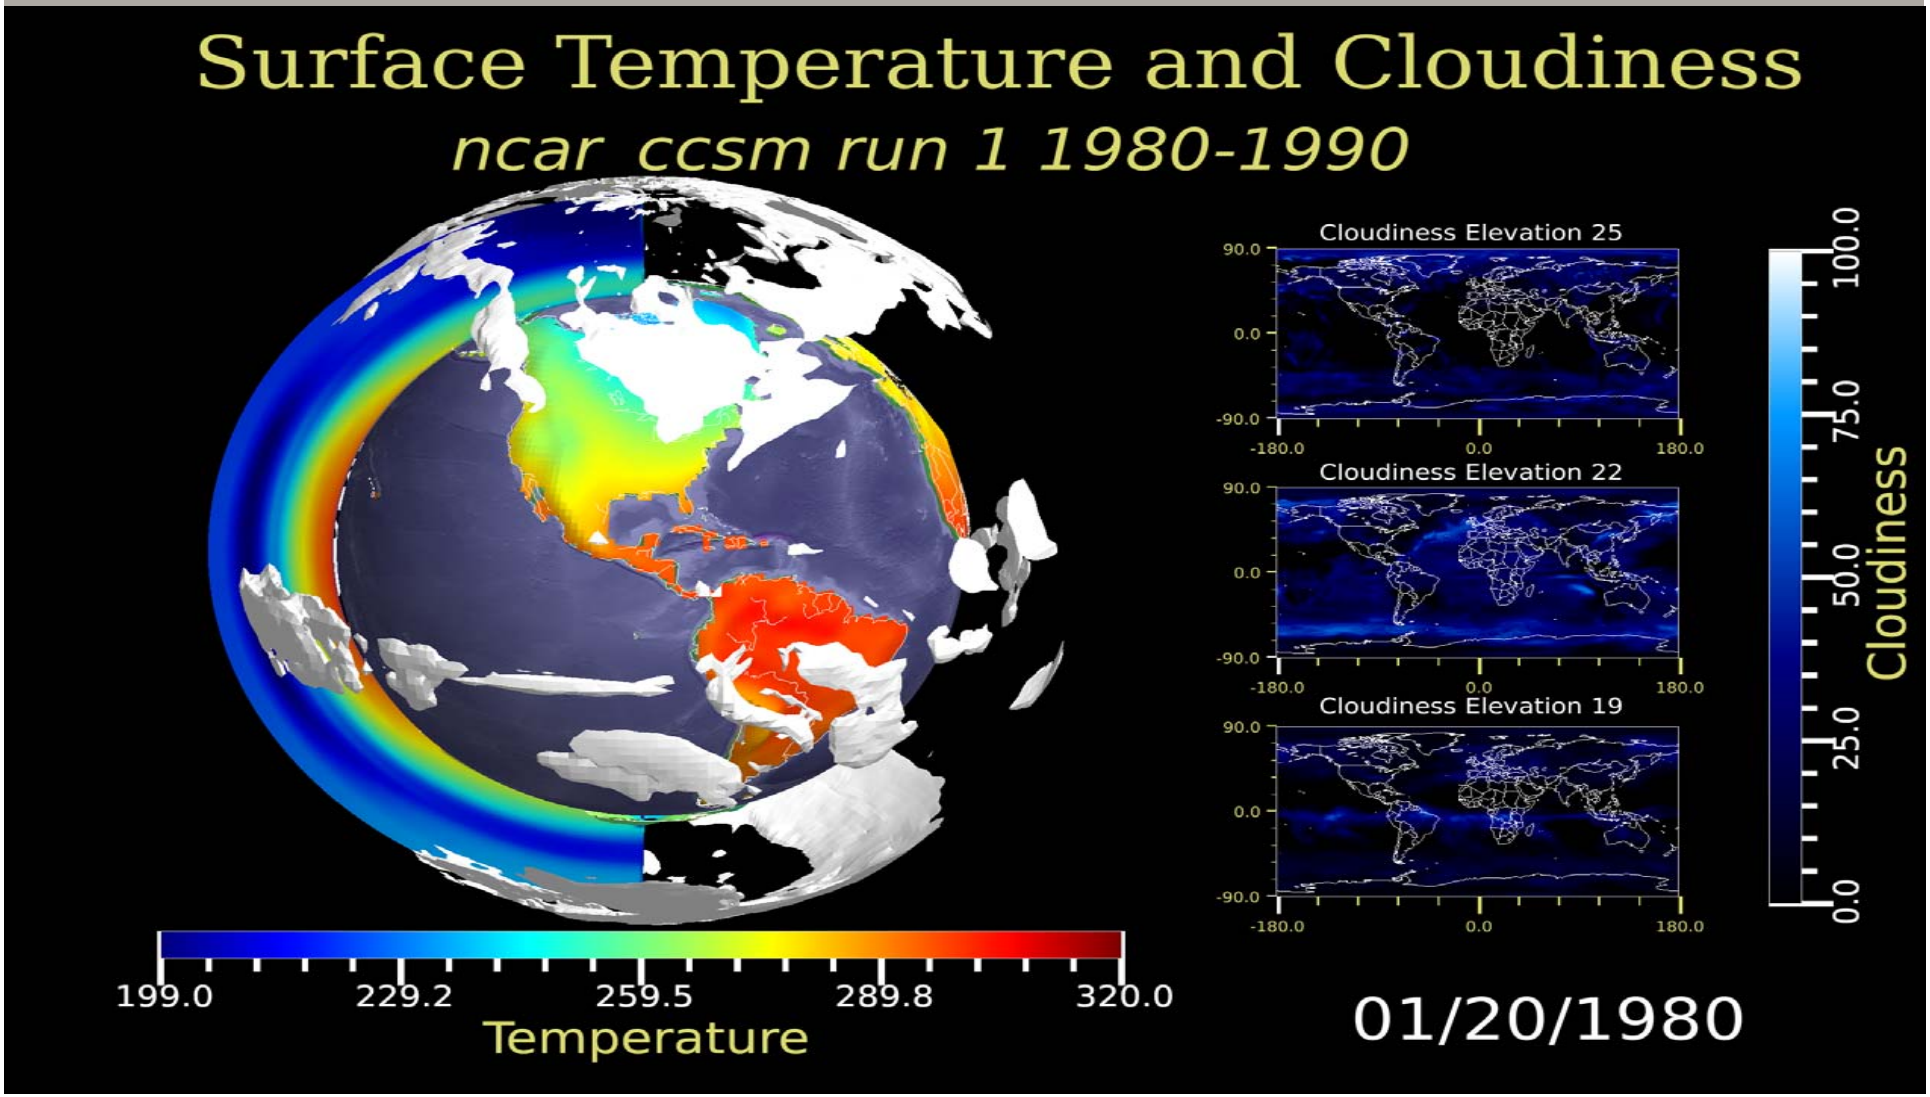
Surface Temperature and Cloudiness
ncar_ccsm run 1 1980-1990
Cloudiness Elevation 25
Cloudiness Elevation 22
Cloudiness Elevation 19
90.0
0.0
-90.0
-180.0
0.0
180.0
100.0
75.0
50.0
25.0
0.0
Cloudiness
199.0
229.2
259.5
289.8
320.0
Temperature
01/20/1980